

Estimating the Energy Use and Efficiency Potential of U.S. Data Centers

These centers house servers, and storage and network devices; this paper provides equations to estimate their total electricity demand and also to estimate potential electricity savings.

By ERIC R. MASANET, RICHARD E. BROWN, ARMAN SHEHABI,
JONATHAN G. KOOMEY, AND BRUCE NORDMAN

ABSTRACT | Data centers are a significant and growing component of electricity demand in the United States. This paper presents a bottom-up model that can be used to estimate total data center electricity demand within a region as well as the potential electricity savings associated with energy efficiency improvements. The model is applied to estimate 2008 U.S. data center electricity demand and the technical potential for electricity savings associated with major measures for IT devices and infrastructure equipment. Results suggest that 2008 demand was approximately 69 billion kilowatt hours (1.8% of 2008 total U.S. electricity sales) and that it may be technically feasible to reduce this demand by up to 80% (to 13 billion kilowatt hours) through aggressive pursuit of energy efficiency measures. Measure-level savings estimates are provided, which shed light on the relative importance of different measures at the national level. Measures applied to servers are found to have the greatest contribution to potential savings.

KEYWORDS | Data centers; energy demand modeling; energy efficiency; information technology

Manuscript received July 10, 2010; accepted January 28, 2011. Date of publication June 30, 2011; date of current version July 20, 2011. This work was supported by the U.S. Environmental Protection Agency, Climate Protection Partnerships Division, Office of Air and Radiation, under U.S. Department of Energy Contract DE-AC02-05CH11231. E. R. Masanet, R. E. Brown, A. Shehabi, and B. Nordman are with the Lawrence Berkeley National Laboratory, Berkeley, CA 94720 USA (e-mail: ermasanet@lbl.gov; rebrown@lbl.gov; ashehabi@lbl.gov; bnordman@lbl.gov).

J. G. Koomey is with the Department of Civil and Environmental Engineering, Stanford University, Stanford, CA 94305 USA (e-mail: jgkoomey@stanford.edu).

Digital Object Identifier: 10.1109/JPROC.2011.2155610

I. INTRODUCTION

As the world shifts from paper-based and analog information systems to digital information management, data centers have become essential to nearly every sector of the global economy. Data centers are facilities that contain information technology (IT) devices used for data processing (servers), storage (storage devices), and communications (network devices). Data centers also contain so-called “infrastructure equipment,” which typically consists of specialized power conversion and backup equipment (to ensure a reliable electricity source), and environmental control equipment (to maintain acceptable temperature and humidity conditions). In the past decade, there has been rapid growth in the number and size of U.S. data centers, with a correspondingly steep rise in electricity demand to power their operations [1]–[3]. The most recent estimates for U.S. data centers suggest that between 2000 and 2006, their electricity demand more than doubled to approximately 61 billion kilowatt hours (kWh) [4], or to around 1.6% of 2006 U.S. electricity sales [5].

The rapid rise and growing national significance of this electricity demand has placed increased attention on strategies for improving the energy efficiency of data center operations [4], [6]–[8]. One prominent example is Public Law 109-431 [9], which in 2007 directed the U.S. Environmental Protection Agency (EPA) to assess U.S. data center electricity demand trends and efficiency opportunities in consultation with a wide audience of IT industry stakeholders. The assessment resulted in a 2007 peer-reviewed report to the U.S. Congress—hereafter referred to as the “EPA study”—containing projections of U.S. data center energy demand under different efficiency

scenarios [4]. The EPA study also contained policy recommendations for promoting greater data center efficiency.

Despite the growing importance of U.S. data center energy use and efficiency, only the EPA study and a handful of other publications comprise the current peer-reviewed, quantitative literature on these topics. These publications are summarized in [3]. Among the most recent of these studies are two by Koomey [1], [2], which presented a bottom-up (i.e., technology-based) model of U.S. data center electricity use based on server installation data from market research firm International Data Corporation (IDC), measured power data by server class, and estimates of infrastructure equipment energy use in 2005. The general bottom-up approach from Koomey [1], [2] was expanded and refined in the 2007 EPA study [4], which further modeled the energy use of storage and network devices within data centers, and also allowed for estimation of energy demand in different data center space types. A novel feature of this model was its ability to estimate the potential electricity savings associated with a select set of broad data center efficiency improvements. The EPA study projected that U.S. data center electricity demand was likely to grow from 61 billion kWh in 2006 to over 107 billion kWh in 2011 in the absence of accelerated efficiency improvements [4]. The study further estimated that 2011 electricity demand could be reduced by as much as 70% through adoption of energy efficient technologies and operating practices.

This paper builds on the previous work by Koomey [1]–[3] and the EPA study [4] in several important ways. First, it documents a concise new mathematical modeling framework for estimating data center energy use and efficiency potentials at different geographic scales, which can be replicated and refined by others. The approach improves the analytical cohesiveness and transparency of the initial model developed for the EPA study, based on extensive feedback from the study’s stakeholder group. The improved model should be accessible to a wider audience, and can be refined as better facility and technology data become available. Second, it provides new insights into the electricity saving potentials and relative importance of specific efficiency measures, whereas the initial model developed in the EPA study only estimated savings associated with broad, non-measure-specific improvements in aggregate fashion. Specifically, the improved model presented here allows one to estimate efficiency potentials associated with discrete efficiency measures applied to different classes of IT devices and infrastructure equipment, and in different space types. Third, this paper applies the improved model to generate the most recent (2008) estimates of both U.S. data center electricity demand and the potential electricity savings associated with nationwide efficiency improvements. These estimates are generated using the most recent available data on the installed base of IT devices and efficiency measures in U.S. data centers.

These estimates should prove more relevant to current research and debates about U.S. data center energy use and efficiency opportunities than previously published estimates.

II. METHODOLOGY

The data center energy model presented here employs a bottom-up modeling approach, which is described in general form by (1). The approach facilitates analysis of energy demand in five data center space types: server closets, server rooms, localized data centers, midtier data centers, and enterprise-class data centers. The characteristics and technology assumptions associated with these data center space types are summarized in Table 1.

This level of spatial disaggregation was chosen because many U.S. servers are expected to be located in server closets and server rooms [10], which have different technology characteristics—and, hence, different efficiency opportunities—than larger data centers. It also facilitates better characterization of electricity costs and potential cost savings, since server closets, server rooms, and localized data centers are often subject to commercial rates whereas larger data centers are often subject to (usually much lower) industrial rates [5]

$$E^{\text{DC}} = \sum_j \left[\sum_i E_{ij}^{\text{S}} + E_j^{\text{ST}} + E_j^{\text{N}} \right] \text{PUE}_j \quad (1)$$

where

E^{DC}	data center electricity demand (kWh/yr);
E_{ij}^{S}	electricity used by servers of class i in space type j (kWh/y);
E_j^{ST}	electricity used by external storage devices in space type j (kWh/y);
E_j^{N}	electricity used by network devices in space type j (kWh/y);
PUE_j	power utilization effectiveness of infrastructure equipment in space type j (kWh/kWh).

Equation (1) estimates data center demand as a function of four variables that account for the electricity use of servers, external storage devices, network devices, and infrastructure equipment. These variables are calculated for each space type using equations and assumptions described in the subsections that follow. In (1), the total electricity use of IT devices within a given space type is determined through summation of the electricity use of servers, external storage devices, and network devices (i.e., the term in brackets). The total electricity use of IT devices is then multiplied by an assumed power utilization effectiveness (PUE) for that space type. The PUE—which is defined as the ratio of total data center energy use to IT device energy use—is a common metric that accounts for the electricity use of infrastructure equipment [11], [12].

Table 1 Typical Characteristics of Data Center Space Types

Space type	Typical size (ft ²)	Typical IT device characteristics	Typical infrastructure equipment characteristics
Server closet	<200	1-2 servers No external storage	Typically conditioned through an office heating, ventilation, and air conditioning (HVAC) system. Environmental conditions are not as tightly maintained as for other data center types. HVAC energy efficiency associated with server closets is probably similar to the efficiency of office HVAC systems.
Server room	<500	A few to dozens of servers No external storage	Typically conditioned through an office HVAC system, with additional cooling capacity, probably in the form of a split system specifically designed to condition the room. The cooling system and backup power equipment are typically of average or low efficiency because there is no economy of scale to make efficient systems more first-cost competitive.
Localized data center	<1,000	Dozens to hundreds of servers Moderate external storage	Typically use under-floor or overhead air distribution systems and a few in-room air conditioning (AC) units. AC units in localized data centers are more likely to be air cooled and have constant-speed fans and are thus relatively low efficiency. Operational staff is likely to be minimal, which makes it likely that equipment orientation and airflow management are not optimized. Air temperature and humidity are tightly monitored. However, power and cooling redundancy may reduce overall system efficiency.
Mid-tier data center	<5,000	Hundreds of servers Extensive external storage	Typically use under-floor air distribution and in-room AC units. The larger size of the center relative to those listed above increases the probability that efficient cooling, e.g., a central chilled water plant and central air handling units with variable speed fans, is used. Staff at this size data center may be aware of equipment orientation and airflow management best practices. However, power and cooling redundancy may reduce overall system efficiency.
Enterprise-class data center	5,000+	Hundreds to thousands of servers Extensive external storage	More efficient equipment is expected to be found in these large data centers. Along with efficient cooling, these data centers may have energy management systems. Equipment orientation and airflow management best practices are most likely implemented. However, enterprise-class data centers are designed with maximum redundancy, which can reduce the benefits gained from operational and technological efficiency measures.

The variables in (1) depend on several parameters related to the adoption of energy efficiency measures as described below. This functionality allows the model to estimate current electricity demand (based on present day adoption of efficiency measures) as well as potential electricity savings in different measure deployment scenarios. The measures included in the model capture the major classes of data center equipment and operations efficiency strategies identified in the EPA study [4], which extensively reviewed such strategies.

An important note is that a number of calculations in the model are made relative to static baseline values that reflect current data center characteristics. This allows estimation of electricity savings potentials between scenarios in a consistent manner. It also reflects a reality in available data; namely, most data on energy saving

measures are expressed relative to current data center practices (e.g., a percent reduction) rather than on an energy intensity basis (e.g., kilowatt hour per computation). Defining energy intensity metrics for data centers is a complex undertaking due to the diversity of services provided; much work is needed before such metrics are available. For clarity, baseline variables in the model are labeled with a “hat” in the remainder of this paper.

A. Servers

Servers are the workhorses of the data center, and as such represent the most significant component (ranging from 50% to over 90%) of IT device electricity demand in all space types [1]–[4]. Correspondingly, servers are the target of numerous efficiency measures. Equation (2) is used to estimate server electricity use by space type based

on server class, the number of servers in each space type, and the annual electricity use per server in each class. The model adopts three server class definitions from IDC based on unit sales prices: volume servers (< \$25 000), midrange servers (from \$25 000 to \$500 000), and high-end servers (> \$500 000). These definitions are used due to the availability of IDC data on U.S. server installations by class [13] and recent power data by class [1]–[4].

Equation (2) estimates the number of installed servers in each class using a baseline value—defined as the current number of installed servers—divided by a “device reduction ratio.” The device reduction ratio accounts for the relative reduction in servers that can occur via efficiency strategies that minimize server counts, such as virtualization, consolidation of applications, and legacy server removal [4]. For example, a device reduction ratio of 3 indicates that three servers have been replaced by one server (i.e., a 3 : 1 reduction ratio). Annual electricity use per server is estimated using (3)–(6), which reflect the relationships between server electricity use and the adoption of key efficiency measures

$$E_{ij}^S = \frac{\hat{N}_{ij}^S}{\rho_{ij}^S} e_{ij}^S \quad (2)$$

where

- E_{ij}^S electricity used by servers of class i in space type j (kWh/y);
- \hat{N}_{ij}^S baseline number of servers of class i installed in space type j ;
- ρ_{ij}^S device reduction ratio for servers of class i in space type j ;
- e_{ij}^S annual electricity use per server of class i in space type j (kWh/y).

Specifically, the potentials for three major efficiency strategies are characterized: 1) use of efficient server hardware; 2) use of dynamic frequency and voltage scaling (DFVS); and 3) reducing the number of physical servers. Efficient server hardware refers broadly to hardware measures such as high-efficiency power supplies, multiple-core processors, more efficient memory, and variable speed fans [4]. Equation (3) expresses the net effect of such measures relative to baseline server electricity use for each server class. DFVS is a common energy saving feature that allows a processor’s clock speed to ramp down during intervals of low utilization, thereby reducing power use. The fractions of a server population with efficient hardware and DFVS enabled can be varied in (3) to estimate server electricity use at different levels of measure adoption

$$e_{ij}^S = \hat{e}_{ij}^S \left(\alpha_{ij}^S \left(\gamma_{ij}^S - 1 \right) + 1 \right) \left(\beta_{ij}^S \delta_{ij}' + \left(1 - \beta_{ij}^S \right) \delta_{ij}'' \right) \quad (3)$$

where

- e_{ij}^S annual electricity use per server of class i in space type j (kWh/y);
- \hat{e}_{ij}^S baseline annual electricity use per server of class i in space type j (kWh/y);
- α_{ij}^S fraction of servers of class i in space type j with energy efficient hardware;
- γ_{ij}^S ratio of efficient server to baseline server electricity use for servers of class i in space type j ;
- β_{ij}^S fraction of servers of class i in space type j with DFVS enabled;
- $\delta_{ij}', \delta_{ij}''$ DFVS and utilization factors.

The net effect of reducing the number of physical servers is captured in (3) through two “DFVS and utilization factors.” These two factors account for the dynamic relationship between the number of installed servers that exist after device reduction initiatives, the average processor utilization of these remaining servers, and the use of DFVS. Fig. 1 plots a representative relationship between server power use, processor utilization, and the state of DVFS (i.e., enabled or disabled) [14]. In virtualization initiatives, several physical servers are replaced by “virtual” servers that reside on a single physical “host” server. An important implication is that the processor utilization of the remaining host servers will rise due to the increased computational demand necessary to support the virtual servers. As is evident in Fig. 1, the rise in processor utilization will lead to an increase in system power use, and the magnitude of this increase depends on the DFVS state (particularly at lower utilization). Despite the increase in server electricity use that accompanies virtualization, data centers can realize substantial electricity savings through large reductions in the number of physical servers.

Equations (4) and (5) calculate the DFVS and utilization factors based on server power-utilization functions such as those illustrated in Fig. 1. For simplicity, and based on available data in [14] and [4], these functions are assumed to be linear and are thus described using slopes and y-axis intercepts in the model (values assumed in this paper are shown in Fig. 1)

$$\delta_{ij}' = \frac{m_{ij}^{\text{ON}} u_{ij} + b_{ij}^{\text{ON}}}{m_{ij}^{\text{OFF}} \hat{u}_{ij} + b_{ij}^{\text{OFF}}} \quad (4)$$

where

- m_{ij}^{ON} slope of power-utilization function (DFVS enabled) for server class i in space type j ;
- u_{ij} postreduction processor utilization per server of class i in space type j (%);
- b_{ij}^{ON} y-intercept of power-utilization function (DFVS enabled) for server class i in space type j ;

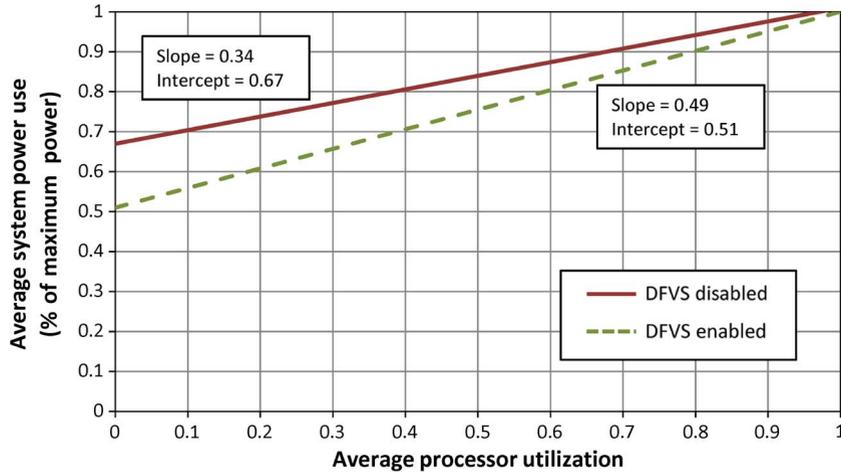


Fig. 1. Relationships between utilization, system power, and DFVS state.

m_{ij}^{OFF} slope of power-utilization function (DFVS disabled) for server class i in space type j ;
 \hat{u}_{ij} baseline processor utilization for active servers of class i in space type j (%);
 b_{ij}^{OFF} y-intercept of power-utilization function (DFVS disabled) for server class i in space type j ;

and

$$\delta_{ij}'' = \frac{m_{ij}^{OFF} u_{ij} + b_{ij}^{OFF}}{m_{ij}^{OFF} \hat{u}_{ij} + b_{ij}^{OFF}} \quad (5)$$

where

m_{ij}^{OFF} slope of power-utilization function (DFVS disabled) for server class i in space type j ;
 u_{ij} postreduction processor utilization per server of class i in space type j (%);
 b_{ij}^{OFF} y-intercept of power-utilization function (DFVS disabled) for server class i in space type j ;
 \hat{u}_{ij} baseline processor utilization for active servers of class i in space type j (%).

The average utilization per server after device reduction is calculated via (6). The postreduction utilization is a function of four variables: 1) the device reduction ratio for servers (defined as the baseline number of installed servers divided by the number that remain after server reduction); 2) the baseline utilization of active servers prior to reduction; 3) the fraction of removed servers that are legacy servers; and 4) the average utilization “overhead” of virtualization software. Legacy servers are those that are functionally obsolete (e.g., hosting applications that are no longer used) but still draw power. Although the presence of legacy servers varies greatly by data center, some industry analysts suggest that they can comprise up to 10% (or more) of the server population at a typical large data

center [22]. For simplicity, it is assumed that legacy servers have negligible utilization and will be completely eliminated in server reduction efforts; thus, they have no effect on postreduction processor utilization. The utilization overhead variable accounts for the processor utilization increase necessary to run virtualization software on the remaining host servers. This software overhead is in addition to utilization increases related to the computational demands of virtual servers

$$u_{ij} = \hat{u}_{ij} \rho_{ij}^S \left(1 - \hat{\theta}_{ij}^S\right) + \hat{u}_{ij} \quad (6)$$

where

u_{ij} postreduction processor utilization per server of class i in space type j (%);
 \hat{u}_{ij} baseline processor utilization for active servers of class i in space type j (%);
 ρ_{ij}^S device reduction ratio for servers of class i in space type j ;
 $\hat{\theta}_{ij}^S$ baseline fraction of servers of class i in space type j that are legacy servers;
 \hat{u}_{ij} postreduction processor utilization overhead per server of class i in space type j (%).

Equations (3)–(6) are designed to assess efficiency opportunities for volume and midrange servers, which account for the vast majority (95%) of U.S. server electricity use [1]–[4]. Efficiency opportunities for high-end servers may be more limited, since they typically incorporate efficient hardware appropriate for their applications (e.g., high-efficiency power supplies) and operate at high utilization (making DFVS less applicable) [4]. However, the approach is equally valid for high-end servers with the appropriate assumptions (see Section III).

B. External Storage

Equation (7) is used to estimate the electricity use of external storage devices by space type. The electricity use of external storage is expressed as a function of the baseline (i.e., current) number of installed devices, the device reduction ratio, baseline storage device electricity use, and assumed adoption levels of key efficiency measures. Equation (7) characterizes the savings potentials associated with two broad efficiency strategies: 1) efficient storage devices and management; and 2) reducing the number of external storage devices. Efficient storage devices and management refers to measures aimed at improving the efficiency of both the physical device [e.g., a switch to high-efficiency hard disk drives (HDDs)] and data management (e.g., tiered storage and/or spinning down HDDs). Device reduction strategies for external storage include such measures as data de-duplication, virtualization, and increasing capacity utilization [4].

Equation (7) can assess any type of external storage device; however, the model currently focuses on external HDD storage and related efficiency opportunities. While tape storage systems are used in many data centers, a lack of data on the installed base and average electricity use of tape storage devices precluded their inclusion in the current model

$$E_j^{ST} = \frac{\hat{N}_j^{ST}}{\rho_j^{ST}} \hat{e}_j^{ST} \left(1 + \alpha_j^{ST} (\gamma_j^{ST} - 1) \right) \quad (7)$$

where

E_j^{ST}	electricity used by external storage devices in space type j (kWh/y);
\hat{N}_i^{ST}	baseline number of external storage devices installed in space type j ;
ρ_j^{ST}	device reduction ratio for external storage in space type j ;
\hat{e}_j^{ST}	baseline annual electricity use per external storage device in space type j (kWh/y);
α_j^{ST}	fraction of energy efficient external storage devices in space type j ;
γ_j^{ST}	ratio of efficient external storage device to baseline external storage device electricity use in space type j .

C. Network Devices

Robust data on the number of installed network devices in U.S. data centers, and their average electricity use, are currently not available in the public domain. Existing reports, audits, and white papers mainly document the relative contribution of network devices to total electricity use at specific facilities [4]. Thus, the model estimates the electricity use of network devices as a fraction of total IT electricity demand for each space type using (8) (rather

than in the bottom-up fashion used for servers and storage devices). In this way, the model enables the use of available (albeit limited) data on network devices in a manner that is consistent with the way those data are reported. Still, the following equation could be used to coarsely estimate the effects of network efficiency improvements by adjusting downward the network device scaling term (i.e., the second term within the brackets)

$$E_j^N = \sum_j \left[\left(\sum_i E_{ij}^S + E_j^{ST} \right) \left(\frac{\varepsilon_j^N}{(1 - \varepsilon_j^N)} \right) \right] \quad (8)$$

where

E_j^N	electricity used by network devices in space type j (kWh/y);
E_{ij}^S	electricity used by servers of class i in space type j (kWh/y);
E_j^{ST}	electricity used by external storage devices in space type j (kWh/y);
ε_j^N	ratio of network device to total IT device electricity use in space type j (kWh/kWh).

D. Infrastructure Equipment

The electricity use of infrastructure equipment is estimated via an assumed PUE for each space type. Equation (9) is used to calculate each PUE, based on assumptions for the electricity use of four major infrastructure system components: power transformers, uninterruptable power supplies (UPSs), cooling systems, and lighting. The cooling systems component represents the broadest class of infrastructure equipment in the model. It refers to primary refrigeration units (e.g., air conditioners and water chillers), coolant pumps, fans and air handlers, cooling towers, and similar equipment. Because the types and configurations of such equipment vary greatly across data centers, cooling system electricity use is represented in aggregate by space type. In (9), the effects of efficiency measures are estimated through changes to the ratio of component to IT device energy demand

$$PUE_j = 1 + \sum_k e_{jk}^I \quad (9)$$

where

PUE_j	PUE of infrastructure equipment in space type j (kWh/kWh);
e_{jk}^I	ratio of electricity use by infrastructure system component k in space type j to IT device electricity use in space type j (kWh/kWh).

Because the PUE is a commonly used metric [9], its use enables the model to leverage reported PUE values from data center audits and benchmarking initiatives [12], [16].

Table 2 Baseline Variable Assumptions

IT Device	Data center space type				
	Server closet	Server room	Localized	Mid-tier	Enterprise
\hat{N}_{ij}^S = Number of installed servers (1,000) [10,13]					
Volume	2,090	2,380	2,040	1,840	3,600
Mid-range	0	18	58	52	240
High-end	0	0	3	2	12
\hat{N}_j^{ST} = Number of installed external storage devices (1,000) [4,17]					
External HDD	0	0	4,390	3,960	8,050
\hat{e}_{ij}^S = Average annual electricity use per server (kWh/y) [1-3]					
Volume	2,060 (for all space types)				
Mid-range	6,910 (for all space types)				
High-end	81,400 (for all space types)				
\hat{u}_{ij} = Average processor utilization (%) [4]					
Volume	10 (for all space types)				
Mid-range	20 (for all space types)				
High-end	70 (for all space types)				
$\hat{\theta}_{ij}^S$ = Fraction of servers that are legacy servers [4]					
Volume	0.05 (for all space types)				
Mid-range	0 (for all space types)				
High-end	0 (for all space types)				
\hat{e}_j^{ST} = Average annual electricity use per external storage device (kWh/y) [4,17]					
External HDD	240 (for all space types)				
Note: Data sources are indicated by bracketed reference numbers.					

However, given its simplistic nature, the PUE is more appropriate for estimating data center energy use in the aggregate than for assessing or comparing the energy use or efficiency of individual facilities [9], [15].

III. SCENARIO AND DATA ASSUMPTIONS

The data center energy model was used to estimate: 1) current (2008) electricity demand of U.S. data centers; and 2) the technically achievable minimum demand assuming maximum adoption of select efficiency measures. These are referred to as the “current demand” and “efficient” scenarios, respectively. The difference between the two scenarios represents the technical potential for electricity savings associated with the selected measures. Technical potentials serve as an upper bound on savings from a technical feasibility perspective; as such, they do not consider factors that may limit the adoption of measures at individual data centers. Such factors could include return on investment criteria, early retirement of existing capital, or perceived risk.

The scenario assumptions are discussed below. All assumptions are based on the best available data in the public domain as of early 2010. For a number of modeling input data in the current demand scenario, the EPA study

[4] remains the most credible (and often only) source of information. As described in Section I, few sources of peer-reviewed data exist in the literature and the EPA study represents the most comprehensive resource for bottom-up, technology-based data among these sources. Furthermore, many of the EPA study data were supplied by the IT industry directly or through industry-led surveys, and all final variable assumptions were subjected to peer review by dozens of IT and data center industry experts. Thus, for many data the EPA study provides reasonable consensus on national average values. Where available, the scenarios employed more recent data as indicated below.

A. Baseline Variables

As discussed in Section II, the model includes static baseline values that reflect current data center characteristics. Assumptions for these variables are summarized in Table 2.

Baseline numbers of installed servers (\hat{N}_{ij}^S) were derived using 2008 market data from IDC [13] and estimated distribution data for server classes across space types previously published in [10]. Based on these data, an estimated total of 12.3 million servers were installed as of 2008. Approximately 97% of these were volume servers and nearly 50% were assumed to be in the largest two space types.

A total 2008 population of 16.4 million external HDDs N_i^{ST} was estimated using market data from IDC [4] and information supplied by a major HDD manufacturer [17]. This total was distributed proportionally across the three largest space types based on installed servers to arrive at the estimates in Table 2.

Baseline IT device electricity use estimates (\hat{e}_{ij}^S and \hat{e}_{ij}^{ST}) were derived from published server [1]–[3] and external HDD [4], [17] power data. Baseline processor utilization (\hat{u}_{ij}) and legacy server fractions $\hat{\theta}_{ij}^S$ were derived from data center survey responses and feedback obtained during the EPA study [4].

B. Scenario Assumptions

Assumptions for the remaining variables are summarized in Table 3, which lists values in the current demand scenario followed by those in the efficient scenario (in parentheses). When a variable does not change between scenarios, only one value is listed; this allows for easy identification of values that change between scenarios, and by how much.

Server device reduction ratios (ρ_{ij}^S) in Table 3 combine the effects of virtualization, application consolidation, and legacy server removal. By default, all server device reduction ratios equal 1 in the current demand scenario. This does not imply that no server reductions have occurred to date; rather, such reductions are already included in the baseline installed server numbers. Adjustments to these ratios in the efficient scenario reflect additional server reductions that could be achieved moving forward. In the efficient scenario, the assumed device reduction ratio for volume servers is 2 for server closets and 5 for all other space types. These values are based on the EPA study [4], which concluded that postreduction server utilization is not likely to exceed 50%–60% in many facilities to ensure a capacity buffer. A device reduction ratio of 2 is assumed for midrange servers in the efficient scenario, since virtualization is increasingly being applied to this server class (with the same assumed capacity buffer constraint as volume servers).

The postreduction utilization overhead per server (\hat{u}_{ij}) equals 0 in the current demand scenario, which assumes that baseline utilization values include existing virtualization software overhead. A value of 10% is assumed in the efficient scenario, commensurate with full deployment of virtualization across the postreduction populations of volume and midrange servers. While virtualization overhead can vary based on software, operating system, and device architecture, a 10% national value was deemed reasonable by stakeholders in the EPA study [4].

Device reduction strategies are not expected to be applicable to high-end servers, given that such servers are expected to operate at high utilization levels [4]. Thus, device reduction ratios and postreduction utilization overhead values for these servers were set to 1 and 0, respectively.

The two scenarios focus on efficient hardware measures for volume servers only. Therefore, no hardware efficiency improvements are assumed between scenarios for midrange and high-end servers; the ratios of efficient server to baseline server electricity use (γ_{ij}^S) equal 1 and the fractions of servers with efficient hardware (α_{ij}^S) equal 0 for both server classes.

For volume servers, the ratio of efficient server to baseline server electricity use equals 0.7. This implies a 30% hardware efficiency improvement, which is based on analyses supporting the recent Energy Star Tier 2 Computer Server Specification [18]. The fraction of volume servers with efficient hardware in the current demand scenario equals 0.05, based on recent market availability of high-efficiency servers from several manufacturers and projections for U.S. sales of such servers in [4]. The efficient scenario assumes that all volume servers have efficient hardware.

DFVS is assumed to be applicable to volume and midrange, but not to high-end, servers, as discussed in Section II. In the current demand scenario, the fraction of servers with DFVS enabled (β_{ij}^S) equals 0.1 for volume and midrange servers, and 0 for high-end servers. These values are based on industry data [23], which suggest that current use of DFVS is quite low despite its widespread availability. The efficient scenario assumes full DFVS use for all volume and midrange servers.

External HDDs are expected to be rare in server closets and server rooms (see Table 1). For these two space types, the device reduction ratios for storage devices (ρ_j^{ST}) equal 1, the ratios of efficient to baseline storage electricity use (γ_j^{ST}) equal 1, and the fraction that is energy efficient (α_j^{ST}) equals 0.

For the other three space types, an achievable HDD reduction ratio of 2 is assumed in the efficient scenario. This value assumes an average capacity utilization of 30%, and that this could be doubled (to 60%) via storage virtualization, data de-duplication, and improved capacity management [24]. The ratio of efficient to baseline storage electricity use equals 0.65 for these three space types in light of two efficiency trends. First, HDD hardware efficiency can be improved through selective adoption of newer high-efficiency HDD technologies (e.g., small form factor HDDs). Second, tiered storage and HDD idling technologies can spin drives down based on data classification and access demands. No robust data exist on the energy savings of these combined strategies; thus, a 35% efficiency improvement was assumed based on data from the EPA study [4], [24]. The current fraction of HDDs operating at this efficiency level was assumed to be low (0.1), based on industry feedback in the EPA study [4]. The efficient scenario assumes that all HDDs will operate at this efficiency level.

The ratio of network device to IT device electricity use (ϵ_j^N) equals 0.05 for server closets, and 0.1 for other space types, based on industry data [25]. These ratios do not

Table 3 Scenario Variable Assumptions

Device/component	Data center space type				
	Server closet	Server room	Localized	Mid-tier	Enterprise
ρ_{ij}^S = Device reduction ratio for servers [4]					
Volume	1 (2)	1 (5)	1 (5)	1 (5)	1 (5)
Mid-range	1	1 (2)	1 (2)	1 (2)	1 (2)
High-end	1	1	1	1	1
\dot{u}_{ij} = Average post-reduction processor utilization overhead per server (%) [4]					
Volume	0 (10)	0 (10)	0 (10)	0 (10)	0 (10)
Mid-range	0 (10)	0 (10)	0 (10)	0 (10)	0 (10)
High-end	0	0	0	0	0
γ_{ij}^S = Ratio of efficient server electricity use to baseline server electricity use [18]					
Volume	0.7	0.7	0.7	0.7	0.7
Mid-range	1	1	1	1	1
High-end	1	1	1	1	1
α_{ij}^S = Fraction of servers with energy efficient hardware [4]					
Volume	0.05 (1)	0.05 (1)	0.05 (1)	0.05 (1)	0.05 (1)
Mid-range	0	0	0	0	0
High-end	0	0	0	0	0
β_{ij}^S = Fraction of servers with DFVS enabled [23]					
Volume	0.1 (1)	0.1 (1)	0.1 (1)	0.1 (1)	0.1 (1)
Mid-range	0.1 (1)	0.1 (1)	0.1 (1)	0.1 (1)	0.1 (1)
High-end	0	0	0	0	0
ρ_j^{ST} = Device reduction ratio for external storage devices [24]					
External HDD	1	1	1 (2)	1 (2)	1 (2)
α_j^{ST} = Fraction of external storage devices that is energy efficient [4]					
External HDD	0	0	0.1 (1)	0.1 (1)	0.1 (1)
γ_j^{ST} = Ratio of efficient storage device electricity use to baseline storage device electricity use					
External HDD	1	1	0.65	0.65	0.65
ϵ_j^N = Ratio of network device electricity use to total IT device electricity use [25]					
Network devices	0.05	0.1	0.1	0.1	0.1
e_{jk}^I = Ratio of infrastructure system component electricity use to IT device electricity use [1-4,12,18,19]					
Transformer	0	0.05 (0.03)	0.05 (0.03)	0.05 (0.03)	0.05 (0.03)
UPS	0	0.20 (0.1)	0.20 (0.1)	0.20 (0.1)	0.20 (0.1)
Cooling	0.95 (0.48)	0.73 (0.36)	0.73 (0.16)	0.73 (0.16)	0.73 (0.16)
Lighting	0.05 (0.11)	0.02 (0.01)	0.02 (0.01)	0.02 (0.01)	0.02 (0.01)

Note: Each cell lists the variable value assumed in the current demand scenario followed by the value assumed in the efficient scenario (in parentheses). When an assumed value does not change between scenarios, only one value is listed. Data sources are indicated by bracketed reference numbers.

change between scenarios based on two simplifying assumptions. First, it is assumed that the number of network ports (and hence network energy use) will decrease proportionally with server counts. It is possible that some data centers would install additional ports on “host” servers to provide additional capacity on network links. However, it is assumed in the efficient scenario that

the number of added ports would be small compared to the number of ports eliminated. Second, network equipment manufacturers and researchers are actively pursuing hardware design and management measures to improve the energy efficiency of network devices (see, for example, [26]). As a preliminary estimate, it was assumed that efficiency gains through such measures could help

Table 4 U.S. Data Center Electricity Use (Billion Kilowatt Hours Per Year) by Space Type

	Data center space type					Total	% of Total
	Server closet	Server room	Localized	Mid-tier	Enterprise		
Current demand (2008) scenario							
Volume servers	4.1	4.7	4.0	3.6	7.1	23.7	34%
Mid-range servers	0.0	0.1	0.4	0.4	1.6	2.5	4%
High-end servers	0.0	0.0	0.2	0.2	1.0	1.4	2%
Storage devices	0.0	0.0	1.0	0.9	1.8	3.7	5%
Network devices	0.2	0.5	0.6	0.6	1.3	3.2	5%
Transformer	0.0	0.3	0.3	0.3	0.6	1.5	2%
UPS	0.0	1.1	1.3	1.1	2.6	6.0	9%
Cooling	4.1	3.9	4.6	4.1	9.4	26.1	38%
Lighting	0.2	0.1	0.1	0.1	0.3	0.8	1%
Total	8.7	10.7	12.6	11.3	25.7	69.0	100%
% of Total	13%	16%	18%	16%	37%	100%	
Efficient scenario							
Volume servers	1.4	0.8	0.7	0.6	1.2	4.6	36%
Mid-range servers	0.0	0.1	0.2	0.2	0.8	1.3	10%
High-end servers	0.0	0.0	0.2	0.2	1.0	1.4	11%
Storage devices	0.0	0.0	0.3	0.3	0.6	1.3	10%
Network devices	0.1	0.1	0.2	0.1	0.4	0.9	7%
Transformer	0.0	0.0	0.0	0.0	0.1	0.2	2%
UPS	0.0	0.1	0.2	0.1	0.4	0.8	6%
Cooling	0.7	0.3	0.3	0.2	0.6	2.2	17%
Lighting	0.2	0.0	0.0	0.0	0.0	0.2	2%
Total	2.3	1.4	2.0	1.8	5.2	12.8	100%
% of Total	18%	11%	16%	14%	41%	100%	
Technical potential for electricity savings							
Volume servers	2.7	3.9	3.4	3.1	6.0	19.1	34%
Mid-range servers	0.0	0.1	0.2	0.2	0.8	1.2	2%
High-end servers	0.0	0.0	0.0	0.0	0.0	0.0	0%
Storage devices	0.0	0.0	0.7	0.6	1.2	2.5	4%
Network devices	0.1	0.4	0.5	0.4	0.9	2.4	4%
Transformer	0.0	0.2	0.3	0.2	0.5	1.3	2%
UPS	0.0	1.0	1.1	1.0	2.2	5.2	9%
Cooling	3.4	3.6	4.3	3.9	8.7	24.0	43%
Lighting	0.1	0.1	0.1	0.1	0.2	0.6	1%
Total	6.4	9.3	10.5	9.5	20.5	56.2	100%
% of Total	11%	17%	19%	17%	36%	100%	

Note: values may not sum to 100% due to rounding.

maintain a constant network device to IT device electricity use ratio in all space types (despite significant reductions in server and storage energy use due to server and storage efficiency improvements). As better bottom-up data emerge on network device energy use and efficiency options, however, these simplifying assumptions should be reassessed.

The ratios of infrastructure system component to IT device electricity use e_{jk}^I correspond to a PUE of 2 for all space types in the current demand scenario. Although PUE values vary widely by facility, a national average of 2 aligns with industry consensus [1]–[4] and available audit data

[12], [18]. All infrastructure component ratios in the current demand scenario were based on the EPA study [4]. In server closets, the PUE is assumed to be a function of two components: building heating, ventilation, and air conditioning (HVAC) systems for IT device heat removal, and lighting (see Table 1).

In the efficient scenario, infrastructure component ratios reflect nationwide average PUE values of 1.6 for server closets, 1.5 for server rooms, and 1.3 for localized, midtier, and enterprise data centers. For server closets and server rooms, average building HVAC efficiency improvements of 50% were deemed feasible based on recent U.S.

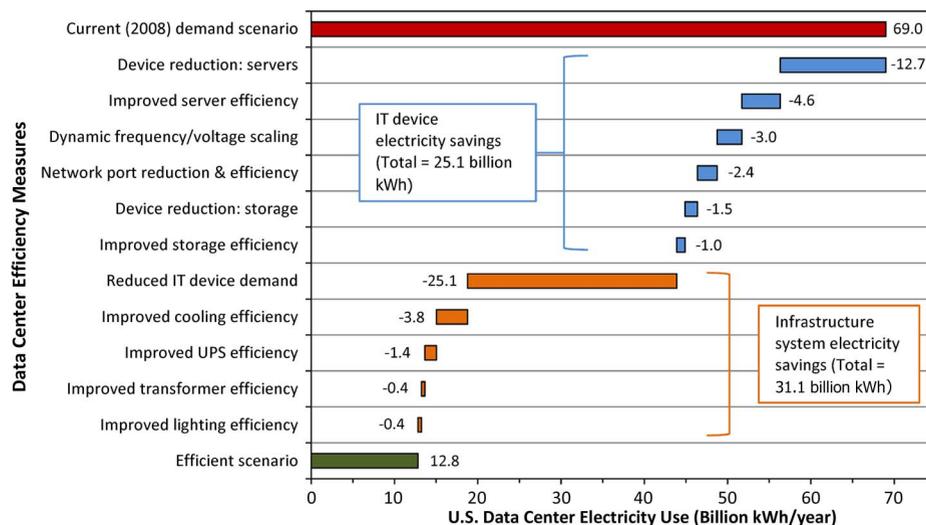


Fig. 2. Efficiency measure contributions to electricity savings.

data for commercial buildings [19]. For localized, midtier, and enterprise data centers, the following improvements were assumed [4]: transformer efficiency improvement from 95% to 98%; UPS efficiency improvement from 80% to 90%; and a shift to cooling best practices (e.g., free cooling, cooling towers, variable-speed air handlers and pumps, and variable-speed drive chillers with economizers). These improvements lead to a nationwide average PUE of 1.3 in these three space types, which aligns well with highly efficient facilities in recent benchmarking studies [12], [16].

Efficiency improvements to transformers and UPS equipment in server rooms were assumed to be similar to those in the larger space types. The component ratios for lighting in the efficient scenario assume that lighting needs are proportional to the number of installed servers, and that lighting efficiency improves by 25% [19].

IV. RESULTS AND DISCUSSION

Table 4 summarizes the results for the current demand and efficient scenarios by IT device, infrastructure system component, and space type. Also provided is a corresponding summary of technical potentials for electricity savings (i.e., the difference between scenario results).

2008 electricity demand is estimated at 69 billion kWh, or around 1.8% of 2008 nationwide electricity sales [5]. This represents a 13% increase from 2006 demand (61 billion kWh) [4]. The increase is largely explained by growth in installed servers, from approximately 11 million in 2006 [4] to over 12.3 million in 2008 [13]. As in previous studies [1]–[4], volume servers and cooling systems are by far the largest components of electricity use; together they accounted for over 70% of current demand. One third of

total demand is estimated to occur in the nation's largest (enterprise) data centers.

Despite continued growth in data center electricity demand, the results for the efficient scenario suggest that deep savings may be achieved through aggressive pursuit of energy efficiency. The technical potential is estimated at approximately 56 billion kWh—an 80% reduction and an amount that is more than double the annual electricity use of Los Angeles (26 billion kWh) [20]. The cost savings from such a reduction would be substantial. Based on 2008 U.S. average electricity rates—10.28 cents/kWh for commercial and 7.01 cents/kWh for industrial buildings—annual electricity costs would be reduced from \$5.9 billion to \$1.1 billion [5]. These results suggest both widespread inefficiencies in current data center operations and the availability of technologies and operating practices that could reduce these inefficiencies significantly. Substantial electricity savings are achievable across all space types, but clearly the largest three (and enterprise data centers in particular) account for the majority of potential savings.

The most significant demand reductions are associated with volume servers and cooling systems, which is expected given their contributions to current data center electricity use (see Table 4). However, Fig. 2 reveals the dominant role that server measures play in reducing electricity use. It plots the average contribution to electricity savings of the efficiency measures assessed by the model. Savings by measure are presented in rank order for IT devices (the top half of Fig. 2, in blue) and infrastructure equipment (the bottom half, in orange). Clearly seen is the dominant role that reduced demand for IT device heat removal and power provision plays in minimizing infrastructure equipment electricity use (indicated in Fig. 2 by “reduced IT device demand”). Of the

31.1 billion kWh infrastructure equipment demand reduction, 25.1 billion kWh are attributable to simply eliminating the need for infrastructure services through reduced IT demand. These results underscore the importance of the well-known “dual benefit” effect of reducing IT device electricity use.¹

Fig. 2 also sheds light on the relative importance of the measures in the current model. Measures for servers offer by far the greatest potential for reducing electricity demand, largely due to device reduction and the adoption of Energy Star compliant volume servers. Considering the “dual benefit” effect, server measures accounted for approximately 70% of the estimated savings. These results underscore the critical importance of efficiency measures for servers in U.S. data centers. Measures for storage devices and the reduction in required network ports accounted for 20% of IT electricity savings.

Although savings for infrastructure equipment are largely attributable to IT device efficiency, Fig. 2 shows that meaningful savings can be realized through infrastructure measures. Improved cooling efficiency is the most important measure, followed by improved UPS efficiency. Electricity savings from transformer and lighting measures are relatively minor, given already high transformer efficiencies and the small contribution of lighting to facility electricity use.

There are several caveats associated with Fig. 2. First, the importance and relative contribution of individual measures will vary by data center, depending on installed equipment, operating practices, space type, location, and other unique factors. Thus, Fig. 2 data should be interpreted only as estimates of national average measure contributions across all data center space types. Second, the relative contribution of measures can change based on the order in which they are applied. Fig. 2 shows the average contribution of each measure based on multiple model runs, which applied measures in different orders. Third, although the relative contribution of infrastructure measures is fairly small, such measures may yield substantial savings in some data centers. Many data centers have significantly reduced electricity demand through such simple improvements as operating at higher temperature set points and improving air flow. However, the relative contribution of infrastructure measures

declines with increasing IT device efficiency due to the “dual benefit” effect.

It is useful to compare results to the EPA study [4]. In its most aggressive “state-of-the-art” technology scenario, the EPA study estimated a nationwide electricity savings potential of around 70%. The somewhat higher estimate of electricity savings in this paper is attributable to two key methodological differences. First, the improved model presented here includes efficiency measures (e.g., mid-range server virtualization, storage efficiency improvements, and Energy Star servers) that were not modeled in the EPA study. Second, the technical potentials presented in this study assume 100% penetration of the stated efficient PUE by space type, whereas the EPA study applied its efficient PUE assumptions to only 50% of data centers within each space type. A penetration of 50% was used in the EPA study as a lower bound on technical potential for infrastructure systems to acknowledge that such improvements may only occur during major equipment upgrades, facility expansions, or new facility construction. Indeed, there are a number of economic, information, and institutional barriers to realizing the full technical potential presented here; such barriers (many are not unique to data centers, and many can be overcome) are discussed in [4], [6], and [27]. Still, the technical potentials presented here are useful for illustrating the full potential of technologies available to data center operators, and for underscoring the extent of the performance gap between technically achievable energy efficiency and real-world practice.

As with any model, the quality and utility of the results depend critically on the availability of credible input data. While the analyses presented here utilized best available data from a wide range of public and industry sources, the robustness of many data could not be verified due to lack of peer-reviewed sources for calibration. Furthermore, a thorough quantitative treatment of uncertainty is not yet possible, given the predominance of point estimates (rather than credible ranges) for many data in the model. Given the bottom-up nature of the model, improved data on installed device numbers, additional device/equipment classes, and device/equipment electricity use in different space types would particularly improve its accuracy. Improved data on tape storage and network devices would further improve the comprehensiveness of the model. Last, the model focuses on electricity use and efficiency. If the use of other fuels becomes more significant (e.g., natural gas engine driven compressors or steam-based absorption chillers) [15], the model can be expanded.

Finally, it is important to understand the macroeconomic context of data center services. Electricity used in data centers enables structural transformations in the economy that can save energy and reduce resource use [28]. For example, a recent analysis comparing the impacts of downloading music to buying it on compact disc (CD) found substantial (40%–80%) savings in carbon emissions for downloads compared to the best case for physical CDs

¹This effect can be visualized via (1) and a simple example. Consider a data center with 100 units of IT device energy demand and a PUE of 2. Equation (1) estimates total data center energy demand of 200 units of energy (100 units for IT devices, 100 units for infrastructure systems). If IT device energy demand is halved (i.e., reduced by 50 units), and the PUE stays constant, total data center energy demand is also halved (i.e., reduced to 100 units: 50 units for IT devices and 50 units for infrastructure systems). Implicit in this effect is the assumption that a data center’s temperature set point remains constant (i.e., reduced heat generation by IT devices will lead to reduced cooling system demand to maintain a constant space temperature).

[21]. Moving bits is often preferable to moving atoms, and while minimizing the direct electricity use of data centers is important, it is also critical to understand the macro-economic system benefits enabled by data centers. ■

Acknowledgment

The authors would like to thank the International Data Corporation (IDC) for data support.

REFERENCES

- [1] J. Koomey, *Estimating Total Power Consumption by Servers in the U.S. and the World*. Oakland, CA: Analytics Press, 2007.
- [2] J. Koomey, *Estimating Regional Power Consumption by Servers: A Technical Note*. Oakland, CA: Analytics Press, 2007.
- [3] J. Koomey, "Worldwide electricity used in data centers," *Environ. Res. Lett.*, vol. 3, art no. 034008, Jul.–Sep. 2008.
- [4] R. Brown, E. Masanet, B. Nordman, W. Tschudi, A. Shehabi, J. Stanley, J. Koomey, D. Sartor, P. Chan, J. Loper, S. Capana, B. Hedman, R. Duff, E. Haines, D. Sass, and A. Fanara, "Report to congress on server and data center energy efficiency: Public law 109-431," Lawrence Berkeley Nat. Lab., Berkeley, CA, LBNL-363E, 2007.
- [5] United States Department of Energy, *Electric Power Annual 2008*, Washington, DC, DOE/EIA-0348(2008), 2010.
- [6] J. Loper and S. Parr, "Energy efficiency in data centers: A new policy frontier, alliance to save energy," Washington, DC, 2007.
- [7] P. Banerjee, C. D. Patel, C. Bash, and P. Ranganathan, "Sustainable data centers: Enabled by supply and demand side management," in *Proc. 46th ACM/IEEE Design Autom. Conf.*, San Francisco, CA, Jul. 26–31, 2009, pp. 884–887.
- [8] A. Rouyer, "Energy efficiency in data centers. Moving forward to face tomorrow's challenges," in *Proc. Int. Conf. Exhib. Power Electron. Intell. Motion Power Quality PCIM Eur.*, Nuremberg, Germany, 2009, ISBN: 978-3-8007-3158-9. [Online]. Available: <http://d-nb.info/993928749/about/html>
- [9] *Code of Federal Regulations, an Act to Study and Promote the Use of Energy Efficient Computer Servers in the United States*, Public Law 109-431, 120 STAT. 2920, 2006.
- [10] M. Bailey, M. Eastwood, T. Grieser, L. Borovick, V. Turner, and R. C. Gray, "Special study: Data center of the future," Int. Data Corp., New York, IDC #06C4799, 2007.
- [11] C. Belady, A. Rawson, J. Pflueger, and T. Cader, "Green grid data center power efficiency metrics: PUE and DCIE," The Green Grid, White Paper #6, 2008.
- [12] S. Greenberg, E. Mills, W. Tschudi, P. Rumsey, and B. Myatt, "Best practices for data centers: Lessons learned from benchmarking 22 data centers," in *Proc. ACEEE Summer Study Energy Efficiency in Buildings*, Asilomar, CA, pp. 3-76–3-87, Aug. 2006.
- [13] International Data Corporation (IDC), *IDC's Worldwide Installed Base Forecast, 2007-2010*, Framingham, MA, 2009.
- [14] *Advanced Micro Devices, Power and Cooling in the Data Center*, Santa Clara, CA, 34246C, 2006.
- [15] K. G. Brill, "Measuring data center energy flow," *Forbes*, Jan. 28, 2009.
- [16] W. Tschudi, T. Xu, D. Sartor, and J. Stein, "High performance data centers: A research roadmap," Lawrence Berkeley Nat. Lab., Berkeley, CA, LBNL-53483, 2003.
- [17] K. Osterberg, "Personal Communication: Data center hard disk drive shipments and power levels," private communication, Mar. 13, 2007, Seagate Technology.
- [18] United States Environmental Protection Agency, *Energy Star Tier 2 Requirements for Computer Servers*, 2009, accessed January 25, 2010.
- [19] R. Brown, S. Borgeson, J. Koomey, and P. Biermayer, "U.S. building-sector energy efficiency potential," Lawrence Berkeley Nat. Lab., Berkeley, CA, LBNL-1096E, 2008.
- [20] Los Angeles Department of Water and Power, LADWP, *Quick Facts and Figures 2009*, 2010, accessed January 25, 2010. [Online]. Available: <http://www.ladwp.com/ladwp/cms/ladwp000509.jsp>
- [21] C. Weber, J. G. Koomey, and H. S. Matthews, "The energy and climate change implications of different music delivery methods," *J. Ind. Ecol.*, vol. 14, no. 5, pp. 754–769, 2009.
- [22] K. G. Brill, "Data center energy efficiency and productivity," presented at the Uptime Institute Symp., Orlando, FL, 2007.
- [23] L. Wigle, "Personal communication," private communication, Jan. 30, 2007, Intel.
- [24] B. Battles, C. Belleville, S. Grabau, and J. Maurier, "Reducing data center power consumption through efficient storage, network appliance," Sunnyvale, CA, Feb. 2007.
- [25] J. Pouchet, "Personal communication," private communication, Mar. 9, 2007, Emerson Network Power, Liebert Corp.
- [26] B. Nordman and S. Lanzisera, "Electronics, networks, and energy," in *Proc. IEEE Conf. Consumer Electron.*, Jan. 2011. [Online]. Available: <http://eetd.lbl.gov/ea/nordman/docs/icce%20nordman%202011.pdf>
- [27] J. G. Koomey, C. Belady, M. Patterson, A. Santos, and K.-D. Lange, *Assessing Trends Over Time in Performance, Costs, and Energy Use for Servers*. Oakland, CA: Analytics Press, 2009.
- [28] E. Masanet and H. S. Matthews, "Exploring environmental applications and benefits of information and communication technology," *J. Ind. Ecol.*, vol. 14, no. 5, pp. 687–691, 2010.

ABOUT THE AUTHORS

Eric R. Masanet received the Ph.D. degree in mechanical engineering from the University of California at Berkeley, Berkeley, with a specialization in environmentally conscious design and manufacturing, in 2004.

He is Deputy Leader of the International Energy Studies Group, Lawrence Berkeley National Laboratory (LBNL), Berkeley, CA. He holds a joint research appointment at the University of California at Berkeley, where he currently serves as Program Manager for the Engineering and Business for Sustainability Certificate Program. He co-led LBNL research efforts (with R. Brown) on the 2007 data center report to the U.S. Congress, with a primary focus on data center efficiency modeling and data analysis. His professional activities include service on the U.S. Technical Advisory Group to the ISO 50001 International Management System Standard for Energy, and as Associate Editor of the journal *Resources, Conservation & Recycling*.



Richard E. Brown received the B.S.E. degree in engineering and management systems from Princeton University, Princeton, NJ, in 1986 and the M.A. degree from the Energy and Resources Group, University of California at Berkeley, Berkeley, in 1993.

He is a Research Scientist in the Environmental Energy Technologies Division, Lawrence Berkeley National Laboratory (LBNL), Berkeley, CA. His current research interests include technical support to the Energy Star program, developing a home energy audit web site (the Home Energy Saver), developing solutions to address the growing energy use of electronics and miscellaneous equipment in buildings, and analyzing the energy use of drinking water and wastewater treatment systems. Before joining the LBNL staff, he spent four years in the Air Force analyzing the cost and performance of satellite systems.



Arman Shehabi received the M.S. degree in environmental engineering from Stanford University, Stanford, CA, in 2000 and the Ph.D. degree in environmental engineering from the University of California at Berkeley, Berkeley, in 2009, with an emphasis in building energy use and indoor air quality.

He is an Energy/Environmental Policy Postdoctoral Fellow at Lawrence Berkeley National Laboratory (LBNL), Berkeley, CA, and has previously held fellowship positions with the National Academy of Science, Washington, DC and with the Consortium on Green Design and Manufacturing, University of California at Berkeley. Between graduate programs, he worked as an engineering consultant and LEED Accredited Professional to develop and implement sustainable building metrics. His current research involves resource and environmental life-cycle assessments of next generation carbon mitigation technologies.



Bruce Nordman received the B.A. degree in architecture and the M.A. energy & resources from the University of California at Berkeley, Berkeley, in 1984 and 1990, respectively.

He is a Research Scientist with the Environmental Energy Technologies Division, Lawrence Berkeley National Laboratory (LBNL), Berkeley, CA, where he has been since 1986. He is a long-time advisor to the U.S. Environmental Protection Agency (EPA) Energy Star program, and active in many technology standards organizations. His work includes topics such as energy use of electronics, low-power mode energy use, network technologies and energy, and user interfaces.



Jonathan G. Koomey received the A.B. in history of science from Harvard University, Cambridge, MA, in 1984, and the M.S. and Ph.D. degrees from the Energy and Resources Group, University of California at Berkeley, Berkeley, in 1986 and 1990, respectively.

He is a Consulting Professor at Stanford University, Stanford, CA. He worked for more than two decades at Lawrence Berkeley National Laboratory (LBNL), Berkeley, CA, and has been a Visiting Professor at Yale University (Fall 2009) and Stanford University (2004–2005 and Fall 2008). He is the author or coauthor of eight books and more than 150 articles and reports, and is one of the leading international experts on the economics of reducing greenhouse gas emissions and the effects of information technology on resource use. His latest solo book is the 2nd edition of *Turning Numbers into Knowledge: Mastering the Art of Problem Solving* (Oakland, CA: Analytics Press, 2008).

